

Our Docket No. 003229.P093
Express Mail No.: EL466329186US

UNITED STATES PATENT APPLICATION

FOR

**DETERMINING NETWORK ADDRESSES AND PORTS USING TABLE
FROM A DESCRIPTION FILE**

INVENTORS:

Earnest E. Hughes
Herbert V. Savage III

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Blvd., 7th Floor
Los Angeles, CA 90025-1026
(714) 557-3800

**DETERMINING NETWORK ADDRESSES AND PORTS USING TABLE
FROM A DESCRIPTION FILE**

RELATED APPLICATION

This application claims the benefit of U.S. Provisional Application No.
5 60/203,138, titled "A System and Method for Determining Network Addresses
and Ports" filed on May 8, 2000.

BACKGROUND

1. Field of the Invention

This invention relates to information processing. In particular, the
10 invention relates to network address translation.

2. Description of Related Art

The need for a Network Address translation arises when a network's
internal addresses cannot be used outside the network either for privacy reasons or
because they are invalid for use outside the network. Basic Address translation
15 would allow hosts in a private network to transparently access the external, or
public, network and enable access to selective local hosts from the outside.
Organizations with a network setup predominantly for internal use, with a need for
occasional external access are good candidates for this scheme. Many Small
Office, Home Office (SOHO) users and telecommuting employees have multiple
20 Network nodes in their office, running applications, but have a single Internet
Protocol (IP) address assigned to their remote access router by their service
provider to access remote networks.

Messages are typically transmitted over the network domain in packets. Packets contain information other than the actual data in the messages. To describe the information and the data in an organized manner, it is useful to have a formal method of abstraction of the packet contents. One such method is the

5 Abstract Syntax Notation One (ASN.1).

ASN.1 is a formal notation used for describing data transmitted by telecommunications protocols, regardless of language implementation and physical representation of these data, whatever the application, whether complex or very simple.

10 One of the main reasons for the success of ASN.1 is that this notation is associated with several standardized encoding rules such as the Basic Encoding Rules (BER), or more recently the Packed Encoding Rules (PER), which prove useful for applications that undergo restrictions in terms of bandwidth. These encoding rules describe how the values defined in ASN.1 can be translated into

15 the bytes 'over the wire' and reverse.

Existing techniques to perform address translation include exhaustively traversing all possible contents for a given data packet to locate addresses and ports. This approach is time consuming and wastes processing resources.

Therefore, there is a need in the technology to provide a technique for

20 network address translation.

SUMMARY

A method and apparatus provides network address translation (NAT) between a first network and a second network. A NAT processor includes a table and a parser. The table stores a plurality of entries for a packet in a sequence of
5 packets of a message transmitted from the first network to the second network. The entries include a first connection identifier corresponding to the first network. The entries are obtained from a description file of the packet. The parser is coupled to the table to parse the sequence of packets using the table. The parser extracts the first connection identifier.

10 According to one embodiment of the present invention, the NAT processor further includes a translator coupled to the parser to translate the first connection identifier into a second connection identifier corresponding to the second network. The table includes at least a message type entry to specify a message type characterizing the message, and at least a data type entry to specify a data type of a
15 subsequent packet in the sequence of packets.

The NAT processor may be able to provide at least one of the following:
(1) efficient processing of packets, (2) compact representation of packet information, and (3) easy applicability to any standard, format, or protocol.

20 Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

5 Figure 1 is a diagram illustrating a system in which one embodiment of the invention can be practiced.

Figure 2 is a diagram illustrating a network address translation (NAT) processing unit according to one embodiment of the invention.

Figure 3 is a diagram illustrating a description file according to one embodiment of the invention.

10 Figure 4 is a flowchart illustrating a process to compile the description file according to one embodiment of the invention.

Figure 5 is a diagram illustrating a data structure for a syntax table according to one embodiment of the invention.

15 Figure 6 is a flowchart illustrating a process to perform network address translation using the syntax table according to one embodiment of the invention.

DESCRIPTION

A method and apparatus provides network address translation (NAT) between a source network and a destination network. The NAT converts a source connection identifier to a destination connection identifier, or vice versa. A
5 connection identifier may be a network address (e.g., IP address), or a port number. A NAT processor includes a table and a parser. The table stores a plurality of entries for a packet in a sequence of packets of a message transmitted from the source network to the destination network. The entries include a first
10 connection identifier corresponding to the source network. The entries are obtained from a description file of the packet. The parser is coupled to the table to parse the sequence of packets using the table. The parser extracts the first connection identifier.

According to one embodiment of the present invention, the NAT processor further includes a translator coupled to the parser to translate the source
15 connection identifier into a destination connection identifier corresponding to the destination network. The table includes at least a message type entry to specify a message type characterizing the message, and at least a data type entry to specify a data type of a subsequent packet in the sequence of packets.

In the following description, for purposes of explanation, numerous details
20 are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention. In other instances, well-known electrical structures and circuits are shown in block diagram form in order not to obscure the present invention. For example, specific details are not
25 provided as to whether the method is implemented in a station as a software routine, hardware circuit, firmware, or a combination thereof.

Embodiments of the invention may be represented as a software product stored on a machine-readable medium (also referred to as a computer-readable

medium, a processor-readable medium, or a computer usable medium having a computer readable program code embodied therein). The machine-readable medium may be any type of magnetic, optical, or electrical storage medium including a diskette, compact disk read only memory (CD-ROM), memory device
5 (volatile or non-volatile), or similar storage mechanism. The machine-readable medium may contain various sets of instructions, code sequences, configuration information, or other data. Those of ordinary skill in the art will appreciate that other instructions and operations necessary to implement the described invention may also be stored on the machine-readable medium. Software running from the
10 machine readable medium may interface with circuitry to perform the described tasks.

Figure 1 is a diagram illustrating a system 100 in which one embodiment of the invention can be practiced. The system 100 includes private networks 110 and 140, a public network 170, and network devices 180 and 190.

15 The private network 110 is a network used within a private entity, such as an organization or an individual home. Examples of the private network 110 include a local area network (LAN) and an intranet. The private network 110 is connected to the public network 170 via the network device 180. The network device 180 may be a network interface, a server, or a gateway. The private
20 network 110 includes K end nodes 120_1 to 120_K , and a router 130. Each of the end nodes 120_1 to 120_K is a device or unit capable of transmitting and/or receiving messages or a sequence of packets 125_1 to 125_K to and/or from the public network 170. Each of the end nodes 120_1 to 120_K may be a computer, a microprocessor, a processing subsystem, a mobile or wireless device, a workstation, a personal
25 digital assistant (PDA) device, etc. The router 130 routes the messages or a sequence of packets 185 between the public network 170 and any one of the end nodes 120_1 to 120_K . The router 130 performs routing function based on quality of service (QoS), traffic condition, network handling capability, and other performance criteria. The router 130 includes a network address translation
30 (NAT) processing unit 135. The NAT processing unit 135 receives a sequence of

packets of a message either from one of the end nodes 120_i to 120_k or from the public network 170 and translates the source network addresses and/or port numbers into a destination network addresses and/or port numbers.

Similarly, the private network 140 may be a LAN or an intranet. The private network 140 is connected to the public network via the network device 190. The network device 190 may be a network interface, a server, or a gateway. The private network 140 includes L end nodes 150_i to 150_L, and a router 160. Each of the end nodes 150_i to 150_L is a device or unit capable of transmitting and/or receiving messages or a sequence of packets 155_i to 155_L to and/or from the public network 170. Each of the end nodes 150_i to 150_L may be a computer, a microprocessor, a processing subsystem, a mobile or wireless device, a workstation, a personal digital assistant (PDA) device, etc. The router 160 routes the messages or a sequence of packets 195 between the public network 170 and any one of the end nodes 150_i to 150_L. The router 160 performs routing function based on quality of service (QoS), traffic condition, network handling capability, and other performance criteria. The router 160 includes a network address translation (NAT) processing unit 165. The NAT processing unit 165 receives a sequence of packets of a message either from one of the end nodes 150_i to 150_L or from the public network 170 and translates the source connection identifier (e.g., network addresses and/or port numbers) into a destination connection identifier (e.g., network addresses and/or port numbers).

The public network 170 is a global network that allows messages to be transmitted and received between the private networks 110 and 140. Other networks may also be connected to the public network 170. For illustrative purposes, only two networks 110 and 140 are shown. Examples of the public network 170 include a wide area network (WAN) and the Internet.

A message may consist of many packets, each packet may include information other than the data contents of the message. This information may include message type, packet/message size, IP address information, and port

identifier. The format of certain data packets may be complicated. A description file is typically generated to describe the packets. An example of such complicated data packets include those packets having different protocols in the audio-video standard employed for video conferencing using, say, NetMeeting®.

- 5 This ASN.1 format description describes the possible contents of a data packet including alternative fields, variable length fields, variable length, optional fields and nested fields. Moreover, this format description includes how such contents are encoded in the data packet.

Figure 2 is a diagram illustrating a network address translation (NAT) processing unit 135/165 according to one embodiment of the invention. The NAT processing units 135 and 165 are essentially the same having similar functionalities. For simplicity, in the following only the reference to the NAT processing unit 135 is used. It is contemplated that the NAT processing unit 165 is similar. The NAT processing unit 135 includes a NAT processor 210, a message description file 220, and a description file compiler 230.

The NAT processor 210 receives the sequence of packets either from one of the end nodes (e.g., 120_1 to 120_K) or the public network 170 and translates the source connection identifiers embedded in the packets into corresponding connection identifiers for the destination. The NAT processor 210 includes a table 240, a parser 250 and a translator 260. The sequence of packets may be the sequence of packets 125, 155, 185, or 195. For simplicity, in the following, reference will be made to the sequence of packets 125. It is understood that other sequence of packets 155, 185, or 195 can be used.

The table 240 stores a plurality of entries for a packet in the sequence of packets 125 of a message transmitted from a source network to a destination network. When one of the end nodes 120_1 to 120_K (or 150_1 to 150_L) sends the message to the public network 170, the source network is the private network 110 (or 140) and the destination network is the public network 170. Similarly, when the public network 170 sends the sequence of packets 185 (or the sequence of

packets 195) to one of the end nodes 120_i to 120_k (or one of the end nodes 150_i to 150_L), the source network is the public network 170 and the destination network is the private network 110 (or the private network 140).

- 5 The entries in the table 240 include a message type entry 242, N data type entries 244₁ to 244_N, and a termination entry 246. One of these entries includes a source connection identifier corresponding to the source network. The message type entry 242 specifies the type of the message (e.g., connect or disconnect messages). Each of the N data type entries 244₁ to 244_N specifies the type of the data in the packets. By examining the type of the data, the parser 250 can
- 10 determine what comes next in the data packets. Each of the N data type entries 244₁ to 244_N may contain a descriptor of the data structure of the packet, such as a field length, a field value, or a pointer to another field. A field in the data type entry may correspond to a network address or a port. The termination entry 246 indicates that there is no more connection identifier in the subsequent packets.
- 15 The termination entry 246 therefore allows the parser 250 to skip the subsequent packets because all network address or port information has been retrieved.

- The parser 250 parses the sequence of packets 124 using the table 240. The parser 250 retrieves the entries from the table 240 and scans the received packets according to the entries. When a source connection identifier such as a
- 20 network address or a port number is encountered, the parser 250 extracts the connection identifier and passes to the translator 260. The translator 260 converts or translates the source connection identifier into the destination connection identifier.

- The entries in the table 240 are generated in advance by the description file
- 25 compiler 230. The description file compiler 230 compiles the description file 220. The description file 220 describes the packets in the message according to a predetermined protocol and/or format. In one embodiment, the format of the description file 220 follows the Abstract Syntax Notation One (ASN.1). As is known by one skilled in the art, any appropriate format or standard can be used.

The table 240, the parser 250, and the translator 260 may be implemented by hardware, software, or a combination of both. When implemented by software, the table 240, the parser 250, and the translator 260 may be represented by modules in a program code. The coupling of these modules is represented by
5 passing parameters, arguments, or pointers. The table 240 may be represented by a data structure.

Figure 3 is a diagram illustrating a description file according to one embodiment of the invention. For illustrative purposes, the description file is written in the Abstract Syntax Notation One (ASN.1). As is known by one skilled
10 in the art, any other syntax or notation can be used.

The notation provides a certain number of pre-defined basic types such as: integers (INTEGER), booleans (BOOLEAN), character strings (IA5String, UniversalString...), bit strings (BIT STRING), etc., and makes it possible to define constructed types such as: structures (SEQUENCE), lists (SEQUENCE
15 OF), choice between types (CHOICE), etc.

Figure 4 is a flowchart illustrating a process 400 to compile the description file according to one embodiment of the invention.

Upon START, the process 400 reads the description file of the packets (Block 410). Then, the process 400 analyzes the syntax rules of the format or
20 standard for the description file (Block 420). Next, the process 400 finds all the possible connection identifiers (e.g., network addresses, port numbers) embedded in the packets (Block 430). Then, the process 400 stores the entries including the message type entries and the data type entries containing the locations of the connection identifiers in the syntax table (Block 440).

25 Then, the process 400 determines if the last connection identifier for the sequence of packets is found (Block 450). If so, the process 400 stores a termination entry in the table indicating that no more addresses or port numbers are contained in the subsequent packets (Block 470) and is then terminated.

Otherwise, the process 400 reads the next module in the description file (Block 460) and returns to Block 430 to continue processing.

Figure 5 is a diagram illustrating a data structure 500 for a syntax table according to one embodiment of the invention.

5 In this illustrative example, the data structure 500 is an array of integer numbers. The numbers correspond to the entries. For example, the entry 510 is the message type entry. Other entries are data type entries which may contain the source connection identifier such as the IP address. Entry 520 is the data type entry which indicates the source network address.

10 Figure 6 is a flowchart illustrating a process 600 to perform network address translation using the syntax table according to one embodiment of the invention.

 Upon START, the process 600 determines if the protocol for the connection is the desired protocol (Block 610). If not, the process 600 processes
15 the packets using the normal protocol (Block 620) and is then terminated. Otherwise, the process 600 retrieves an entry from the table (Block 630). Next, the process 600 determines if the entry contains or points to a connection identifier (Block 640). If not, the process 600 goes to Block 680. Otherwise, the process
20 600 obtains the connection identifier from the packet (Block 650). Next, the process 600 translates the source network address and/or the source port number into the destination network address and/or the destination port number (Block 660). Then, the process 600 obtains the next entry in the table (Block 670). Next, the process 600 returns to Block 640.

 In Block 680, the process 600 determines if the entry is a termination
25 entry. If so, the process 600 is terminated. Otherwise, the process 600 parses the packet based on the entry (Block 690) and then goes to block 670 to continue retrieving entries from the table.

While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to
5 which the invention pertains are deemed to lie within the spirit and scope of the invention.

003239.P093